



ERDC MSRC PET Technical Report No. 01-14

**Improvements in Parallel Chimera
Grid Assembly**

by

Nathan C. Prewitt
Davy M. Belk
Wei Shyy

15 May 2001

**Work funded by the Department of Defense
High Performance Computing Modernization Program
U.S. Army Engineer Research and Development Center
Major Shared Resource Center through**

Programming Environment and Training

Supported by Contract Number: DAHC94-96-C0002
Computer Sciences Corporation

Views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of Defense position, policy, or decision unless so designated by other official documentation.

Improvements In Parallel Chimera Grid Assembly

Nathan C. Prewitt*

Davy M. Belk†

Wei Shyy‡

1 Introduction

The Chimera grid scheme [1] (i.e., the use of overlapping, structured grids for domain decomposition) allows the calculation of moving-body, fluid dynamics problems using time-accurate computational fluid dynamics (CFD). Thus, Chimera has seen application in the analysis of store separation [2], tilt-rotor aircraft aerodynamics [3], and many other areas. However, an expensive grid assembly process, involving hole cutting and interpolation, is required to establish communication between the overlapping grids. And the expense increases further with the time-accurate solution of moving-body problems, because the grid assembly has to be updated after each time-step.

There have been numerous attempts to create robust and efficient algorithms for handling the grid assembly problem. The number of algorithms that have been applied to dynamic, moving-body problems is somewhat less. The number of parallel implementations for dynamic problems is fewer still.

Meakin et al. [4, 5] have performed the parallel solution of many dynamic problems using Chimera methods. However, their decomposition of the work related to the grid assembly function was based on the same decomposition used for balancing the work of the flow solver. Modifications in the work distribution of the grid assembly lessened the parallel performance of the flow solver, and thus, the parallel performance of the overall simulation [5].

Prewitt et al. [6] performed the first parallel, Chimera CFD solution of a dynamic problem that used a data decomposition of the grid assembly separate from that of the flow solver. However, large load imbalances were seen in the grid assembly function, and the relatively small number of grids used in the test problem did not offer much opportunity for scalability. This work presents the improvements in load balancing offered by using a fine-grain decomposition of the work associated with grid assembly.

2 Grid Assembly

There are two major steps in the grid assembly function: hole cutting and the stencil search. Holes are cut into overlapping grids to mark cells that intersect solid boundaries (rather

*PET Onsite CFD Lead, Mississippi State University, U.S. Army ERDC MSRC

†Technical Advisor, Computational Mechanics Branch, Assessments and Demonstrations Division, Air Force Research Laboratory

‡Professor and Chairman, Department of Aerospace Engineering, Mechanics and Engineering Science, University of Florida

than conforming to them) as invalid. This creates hole boundaries that require some sort of applied boundary condition. The points on the fringe of these hole boundaries and the points on the outer boundaries of any embedded grids are collectively referred to as intergrid boundary points (IGBPs). Boundary conditions are supplied at the IGBPs by interpolating the flow solution from any overlapping grids. Finding appropriate interpolation stencils is the goal of the stencil search.

In this work, the concern is with the grid assembly algorithms used in the Beggar code [7]. Beggar is capable of solving three-dimensional inviscid and viscous flow problems involving multiple moving objects. It allows blocked, patched, and overlapping structured grids in a framework that includes grid assembly, flow solution, force and moment calculation, and the integration of the rigid body, six degrees-of-freedom (6DOF) equations of motion.

For cutting holes, Beggar uses an outline and fill algorithm. The facets of the hole-cutting surfaces are used to create an outline of the hole, and then the hole is filled by a fast sweep through the grid. The algorithm first marks the cells surrounding the vertices of a hole-cutting facet as being on the hole side or the world side of the facet. Then the facet is recursively refined in order to ensure a complete outline of the hole. Each of the hole-cutting facets is treated independently. The only restriction on the hole-cutting facets is that they form completely closed surfaces.

3 Parallel Implementations

Beggar has gone through an evolutionary process of parallel-processing development. Single program, multiple data (SPMD) and multiple program, multiple data (MPMD) programming models have been used along with the message-passing paradigm. Implementations have used both PVM and MPI for the message-passing interface.

Belk and Strasburg [8] developed the first parallel implementation. The grid assembly was performed as an initialization step, and only the flow solver was executed in parallel. Thus, this implementation was only applicable to static problems.

The extension to solving dynamic problems in parallel was accomplished in several phases based on the parallel implementation of the grid assembly function and its integration with the parallel implementation of the flow solver. Prewitt et al. [9] presented the first of these implementations, Phase I, which used a single front-end (FE) process to perform the grid assembly in a serial fashion with respect to the parallel execution of the flow solver across multiple back-end (BE) processes. Proper communication between the flow solver and the grid assembly function was required; however, no consideration was given to load balancing or partitioning of the grid assembly function.

In the Phase II implementation, Prewitt et al. [9] showed an improvement in the parallel efficiency by hiding the execution time of the grid assembly function behind that of the flow solver. This was possible because of the Newton-Relaxation scheme used in the flow solver (complete details on the flow solver can be found in the literature; see Whitfield [10], for example).

As the number of BE processors increases, the time to compute the flow solution, and therefore the time available for hiding the work of the grid assembly, decreases. In order to continue to see the optimum speedup, multiple processors must be used to decrease the total execution time of the grid assembly function.

In the Phase III implementation, Prewitt et al. [6] distributed the work of the grid assembly across multiple FE processes using data decomposition techniques. The basis for decomposition was the grids or groupings of grids called *superblocks*. The grid assembly execution times were monitored, and the superblocks were migrated between FE processes in order to dynamically improve the balance of the grid assembly workload.

This implementation showed significant improvements over the previous implementations; however, its performance was hampered by a poor load balance and limited scalability in the grid assembly function, because of the use of superblocks as the basis for domain decomposition. The load balancing and scalability can be improved by using a fine-grain decomposition of the work associated with the grid assembly.

The work of the grid assembly function is associated with the number of hole-cutting surface facets, the number of cells that are cut, and the number of cells that require interpolation. In this article, the hole-cutting facets are used as the basis for decomposition of the hole cutting. This is done to demonstrate the improvements in performance and scalability that are possible with a fine-grain decomposition. Of course, hole cutting is only half of the work of the grid assembly. The other half, the stencil search, is more difficult to implement using a fine-grain decomposition (e.g., based on IGBPs); therefore, the stencil search portion of the grid assembly function is still decomposed based on the mapping of superblocks to FE processes.

4 Parallel Performance

If one considers the entire workload of a complete simulation to be broken down into a part that can be executed in parallel and a part that must be executed serially, the maximum speedup that can be achieved is determined by Amdahl's law [11]:

$$s \leq \frac{1}{\left(\frac{f_p}{n_{pes}}\right) + f_s} \quad (1)$$

where s is the speedup, f_s is the serial fraction of the work, f_p is the parallel fraction of the work, and n_{pes} is the number of processors on which the parallel portion of the code is running.

Generalizing Equation 1 to include the effect of a load imbalance, the speedup becomes

$$s \simeq \frac{1}{\left(\frac{f_p}{n_{pes}}\right) * f_i + f_s} \quad (2)$$

where f_i is the load imbalance factor. This is a scale factor that is used to increase the average parallel work per processor to represent the maximum work on any processor.

The Phase I implementation fits this model of performance very well. The flow solver represents the parallel work, while the grid assembly represents the serial work. Thus, for the Phase I implementation, the speedup can be approximated from Equation 2 as:

$$s \simeq \frac{1}{\left(\frac{F_p}{n_{bes}}\right) * F_i + G_s} \quad (3)$$

where F_p is the parallel fraction of the work represented by the flow solver, F_i is the load imbalance factor from the distribution of the flow solver, $nbes$ is the number of BE processes, and G_s is the serial fraction of the work represented by the grid assembly.

For the Phase II implementation, one must take into account the fact that the execution time of the grid assembly may be hidden by the execution time of the flow solver. If the fraction of time spent in the flow solver after the first Newton step, or *dt iteration*, is

$$F_t = \left(\frac{F_p}{nbes} \right) * F_i * \left(\frac{ndt - 1}{ndt} \right) \quad (4)$$

where ndt is the number of dt iterations being run per time-step, then the speedup can be approximated by the equation

$$s \simeq \frac{1}{\left(\frac{F_p}{nbes} \right) * F_i + G_t} \quad (5)$$

where

$$G_t = \begin{cases} G_s - F_t & \text{if } F_t < G_s \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

If F_t is greater than G_s , the time to do the grid assembly is completely hidden by the flow solver. If the time to compute the grid assembly is only partially hidden by the flow solver, the speedup is degraded by the portion of the grid assembly process that is not hidden.

For the Phase III implementation, the grid assembly is executed in parallel; therefore, the grid assembly time is

$$\frac{G_p}{nfes} * G_i \quad (7)$$

where G_p is the parallel fraction of the work represented by the grid assembly function, G_i is the load imbalance in distributing the grid assembly work, and $nfes$ is the number of FE processors that are executing the grid assembly function. Functional overlapping of the flow solver and the grid assembly function is still being used; therefore, the speedup is still estimated using Equation 5 with

$$G_t = \begin{cases} \frac{G_p}{nfes} * G_i - F_t & \text{if } F_t < \frac{G_p}{nfes} * G_i \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

With the change in the basis for the decomposition of the hole-cutting work, the grid assembly time should be broken down into hole-cutting time and stencil-searching time for the purpose of performance modeling so that separate load imbalance factors can be identified. However, equations 5 and 8 will continue to be used. Any improvement in the overall performance is due to improvements in the load balance of the hole cutting.

5 Results

The generic store, ripple release case is being used for collecting timing information. Three stores are ejected from a triple ejector rack (TER) configuration in bottom, outboard, inboard order with 0.04 seconds between releases. A sense of the motion and the overlapping




Figure 1: Three store, ripple release overlapping grid system with holes

grid system can be obtained from Figure 1. The complete calculation consists of 600 time-steps of 0.0005 seconds, representing 0.3 seconds of the trajectory. Two dt iterations and six inner iterations are used per time-step of the flow solver. This problem was first presented by Thoms and Jordan [12]. Complete details of the problem definition can be found in that reference.

All solutions were run in double precision (64 bit) on an SGI Origin 2000 machine. This particular machine was configured with 64 - 195 MHz R10000 processors and 16 GB of distributed, shared memory with two processors and 512 MB of memory per node card.

A single-processor simulation was computed to establish the base solution time of 9,384 minutes (about 6.5 days). Comparing the execution time of the grid assembly function to the total execution time from the sequential run, work fractions of $G_p = .05$ and $F_p = .95$ were established.

All of the Phase III runs used 4 FE processes for grid assembly, with dynamic load balancing performed after each iteration. Performance data, originally presented by Prewitt et al. [6], are shown in Figure 2. Two estimated speedup curves are shown. One is for nominal load imbalance factors of $F_i = 1.05$ and $G_i = 1.08$. This represents the level of performance that was anticipated. However, the grid assembly experienced a much larger load imbalance, and the overall performance suffered. The second estimated speedup curve shows that the actual imbalance in the grid assembly is about 65 percent.

The iterative algorithm used by the grid assembly function requires several synchronization points to ensure that each process has access to the proper cell state information (i.e.,




Figure 2: Measured speedups for the Phase III runs

iblanking information). The large imbalance is caused by this need for synchronization and the limited number of superblocks with which to decompose the workload.

Figure 3 shows the evolution of the execution times for hole cutting from the 40+4 (40 BE and 4 FE processes) Phase III run. Each curve represents the execution time versus the time-step number for the work done on an FE process. If the workload was well balanced, the FE processes would be executing for nearly equal times. However, the separation in these curves shows that considerable differences exist in the execution time of each process at any given time-step.

There are about 60,000 hole-cutting facets and only 10 superblocks in this particular test case. Therefore, the use of the hole-cutting facets as a basis for decomposition adds flexibility for distributing the workload. The use of shared memory to store the cell state information also reduces communication and the need for synchronization.

Figure 4 shows plots of the hole-cutting execution times versus the time step number when the hole-cutting facets were used as the basis for decomposition of the hole-cutting workload. The upper plot represents a run on 4 FE processes while the lower plot represents a run on 8 FE processes. Although it is difficult to tell, there is a separate curve for each of the FE processes. The tight grouping of the curves shows the excellent load balance experienced by the hole cutting function. A direct comparison of the two plots shows that the execution time decreased by a factor of 2. This, again, is ideal behavior, since the number of FE processes was doubled.




Figure 3: Detailed execution times for hole cutting from the 40+4 Phase III run

During the initial time-steps and at several points during the simulation, there appears to be noise in the timing data. These are variations in the execution times of different FE processes indicating some load imbalance. However, the dynamic load balancing routine quickly drives the execution times back toward the ideal load balance.

Figure 5 shows the speedup for the entire simulation when a fine-grain decomposition of the hole-cutting workload is used along with dynamic load balancing. The additional data are for runs on 4 FE processes and 28, 32, 36, and 40 BE processes. Only slight improvements in the overall speedup are seen since only the load imbalance of the hole cutting is being improved. A fine-grain decomposition of the stencil search must be included in order to reap the full benefits.

6 Conclusion

Execution timing data for the ripple release test problem were presented to show the performance improvements seen by using a fine-grain decomposition of the work associated with hole cutting, as opposed to a coarse-grain decomposition based on the superblocks. This implementation used shared memory to store the cell state information, thereby reducing the communication cost and the requirements for synchronization. A dynamic load balancing algorithm was also used to keep the load balance of the hole-cutting work near ideal.

Several equations were derived from Amdahl's law for use in modeling the performance of the different parallel implementations of the grid assembly function. These equations proved effective in approximating the speedup to be gained and are of great utility in deciding the optimum number of processors on which to run a job.

7 Acknowledgments

This work was performed in part under the funding and oversight of the Air Force Seek Eagle Office. Support was provided in part by a grant of computer time from the DoD High Performance Computing Modernization Program at the U.S. Army Space and Strategic Defense Command Distributed Center, Aeronautical Systems Center Major Shared Resource Center, and U.S. Army Engineer Research and Development Center Major Shared Resource Center. Additional support was provided by the Air Force Office of Scientific Research, Air Force Materiel Command, USAF, under grant number F499620-98-1-0092. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon.




Figure 4: Measured execution times for hole cutting on 4 (top) and 8 (bottom) FE processes with a decomposition based on hole-cutting facets.




Figure 5: Comparison of speedup including fine-grain hole cutting with dynamic load balancing.

References

- [1] Benek, J.A., Steger, J.L., and Dougherty, F.C., "A Flexible Grid Embedding Technique with Applications to the Euler Equations," AIAA Paper 83-1944, June 1983.
- [2] Lijewski, L.E., and Suhs, N., "Time-Accurate Computational Fluid Dynamics Approach to Transonic Store Separation Trajectory Prediction," *Journal of Aircraft*, Vol. 31, No. 4, pp. 886-891, August 1994.
- [3] Meakin, R.L., "Moving Body Overset Grid Methods for Complete Aircraft Tiltrotor Simulations," AIAA Paper 93-3350-CP, AIAA Computational Fluid Dynamics Conference, July 1993, pp. 576-588.
- [4] Barszcz, E., Weeratunga, S.K., and Meakin, R.L., "Dynamic Overset Grid Communication on Distributed Memory Parallel Processors," AIAA Paper 93-3311, July 1993.
- [5] Wissink, A.M., and Meakin, R.L., "On Parallel Implementations of Dynamic Overset Grid Methods," SC97: High Performance Networking and Computing, November 1997.
- [6] Prewitt, N.C., Belk, D.M., and Shyy, W., "Distribution of Work and Data for Parallel Grid Assembly," AIAA Paper 99-0913, January 1999.
- [7] Maple, R.C., and Belk, D.M., "Automated Set Up of Blocked, Patched, and Embedded Grids in the Beggar Flow Solver," *Numerical Grid Generation in Computational Fluid Dynamics and Related Fields*, ed. N.P. Weatherill et al., Pine Ridge Press, pp. 305-314, 1994.
- [8] Belk, D.M., and Strasburg, D.W., "Parallel Flow Solution on the T3D with Blocked and Overset Grids," 3rd Symposium on Overset Composite Grid and Solution Technology, November 1996.
- [9] Prewitt, N.C., Belk, D.M., and Shyy, Wei, "Implementations of Parallel Grid Assembly for Moving Body Problems," AIAA Paper 98-4344, August 1998.
- [10] Whitfield, D.L., "Newton-Relaxation Schemes for Nonlinear Hyperbolic Systems," Engineering and Industrial Research Station Report MSSU-EIRS-ASE-90-3, Mississippi State University, October 1990.
- [11] Amdahl, G.M., "Validity of the Single-Processor Approach to Achieving Large Scale Computing Capabilities," AFIPS Conference Proceedings, Vol. 30, pp. 483-485, April 1967.
- [12] Thoms, R.D., and Jordan, J.K., "Investigations of Multiple Body Trajectory Prediction Using Time Accurate Computational Fluid Dynamics," AIAA Paper 95-1870, June 1995.